arXiv:1510.07244v1 [cs.MS] 25 Oct 2015

The efficiency of boundary element methods depends crucially on the time required for setting up the stiffness matrix. The far-field part of the matrix can be approximated by compression schemes like the fast multipole method or $\mathcal{H}$-matrix techniques. The near-field part is typically approximated by special quadrature rules like the Sauter-Schwab technique that can handle the singular integrals appearing in the diagonal and near-diagonal matrix elements.

Since computing one element of the matrix requires only a small amount of data but a fairly large number of operations, we propose to use GPUs to handle vectorizable portions of the computation: near-field computations are ideally suited for vectorization and can therefore be handled very well by GPUs. Modern far-field compression schemes can be split into a small adaptive portion that exhibits divergent control flows and is handled by the CPU and a vectorizable portion that can again be sent to GPUs.

We propose a hybrid algorithm that splits the computation into tasks for CPUs and GPUs. Our method presented in this article is able to speedup the setup time of boundary integral operators by a significant factor of 19–30 for both the Laplace and the Helmholtz equation in 3D when using two consumer GPGPUs compared to a quad-core CPU.

Approximation of boundary element matrices using GPGPUs

Steffen Börm and Sven Christophersen

October 27, 2015

1 Introduction

Compared to finite element methods, boundary element methods offer advantages like improved accuracy and the possibility to handle infinite domains, but these advantages come at the price of a stiffness matrix that is no longer sparse and that requires us to evaluate singular integrals.

During the last decades, both issues have been the focus of significant research activities: approximation schemes like the fast multipole method [25, 15, 1, 13, 32, 24], panel clustering [19, 27], or hierarchical matrices [16, 17, 2, 14, 7] replace the stiffness matrix by a data-sparse approximation of arbitrary accuracy that can be computed in almost linear complexity. State-of-the-art compression schemes relying on $\mathcal{H}^2$-matrices [18, 8, 4, 3, 5] can handle matrices with millions of degrees of freedom in minutes on modern computers.

There have been some attempts in the past to utilize the computational power of GPGPUs to reduce setup times for discretized boundary integral operators [31, 23, 20, 30], but as far as we know none of them work in the framework of $\mathcal{H}$- and $\mathcal{H}^2$-matrices yet.

The singular integrals that arise in the setup of boundary integral operators can be evaluated directly in special cases, but frequently quadrature rules [29, 26, 12, 28] are employed because they offer more flexibility and can also be more efficient. In modern implementations of boundary element methods, a large portion of the total runtime is due to the evaluation of the singular and near-singular integrals. Our goal is to improve the performance of BEM implementations by using general-purpose graphics processors (GPGPUs) to speed up the computation of these integrals.

Typical GPGPUs differ from standard processors in several important aspects:

- they can contain hundreds or thousands of processing elements, and this allows them to reach a theoretical peak performance in the teraflop range, but
- the processing elements are not independent, instead they are organized in groups that act like vector processors, i.e., in each cycle, all processing elements carry out the same operation (SIMD), and

- although the memory bandwidth is significantly larger than for standard processors, it is still fairly low compared to the large number of processing elements requiring input.

Fortunately, the evaluation of quadrature rules is a task that requires a large number of operations, but only a small amount of data, so the relatively low memory bandwidth does not have a significant impact.

Certain quadrature rules [12, 28] sort integrals into one of a small number of categories and then perform identical operations with each member of these categories. Heterogeneous systems are a good match for these rules: the main processor classifies the integrals according to the categories and sorts them into work packages for the GPGPUs. These packages are ideally suited for vector processors, therefore they are sent to the GPGPUs that evaluate the quadrature rule for each integral. Once all integrals in a work package have been handled, the results are returned to the main processor to be added to the final stiffness matrix.

Even with comparably inexpensive mainstream GPGPUs, this approach can drastically reduce the runtime for singular and near-singular integrals.

Taking advantage of the recently developed Green cross approximation (GCA) method [6, 5], we can use a very similar approach to also construct an efficient $\mathcal{H}^2$-matrix approximation of the remainder of the matrix.

2 Sequential algorithms

We consider the construction of an $\mathcal{H}^2$-matrix approximation of the stiffness matrix of a Galerkin BEM discretization. Given finite element bases $(\varphi_i)_{i\in\mathcal{I}}$ and $(\psi_j)_{j\in\mathcal{J}}$, a kernel function $g : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{C}$, and a sub-manifold $\Omega \subseteq \mathbb{R}^d$, the entries of the stiffness matrix $G \in \mathbb{C}^{\mathcal{I}\times\mathcal{J}}$ are given by

$$g_{ij} = \int_\Omega \varphi_i(x) \int_\Omega g(x,y)\psi_j(y)\,dy\,dx \qquad \text{for all } i \in \mathcal{I},\ j \in \mathcal{J}. \tag{1}$$

Following the standard paradigm of finite element methods, the set Ω is represented by a triangulation $\mathcal{T}$ consisting of triangles $\Delta \in \mathcal{T}$. Representing the entries of G by

$$g_{ij} = \sum_{\substack{\Delta_\varphi\in\mathcal{T}\\ \varphi_i|_{\Delta_\varphi}\neq 0}} \sum_{\substack{\Delta_\psi\in\mathcal{T}\\ \psi_j|_{\Delta_\psi}\neq 0}} \int_{\Delta_\varphi} \varphi_i(x) \int_{\Delta_\psi} g(x,y)\psi_j(y)\,dy\,dx \qquad \text{for all } i \in \mathcal{I},\ j \in \mathcal{J},$$

the task of computing these entries reduces to the task of evaluating integrals on the few pairs $(\Delta_\varphi, \Delta_\psi)$ of triangles where φ_i and ψ_j do not vanish.

Assume that $\Delta_\varphi, \Delta_\psi \in \mathcal{T}$ are given. In order to evaluate the integral

$$\int_{\Delta_\varphi} \int_{\Delta_\psi} \varphi_i(x) g(x,y) \psi_j(y)\,dy\,dx,$$

we introduce the reference triangle

$$\widehat{\Delta} := \{(s,t) \in \mathbb{R}^2 \ : \ 0 \leq t \leq s \leq 1\}$$

and affine maps $\Phi_\varphi : \widehat{\Delta} \to \Delta_\varphi$ and $\Phi_\psi : \widehat{\Delta} \to \Delta_\psi$. By the definition of surface integrals, we only have to evaluate

$$\sqrt{\det(D\Phi_\varphi^T D\Phi_\varphi)}\sqrt{\det(D\Phi_\psi^T D\Phi_\psi)} \int_{\widehat{\Delta}} \int_{\widehat{\Delta}} \varphi_i(\Phi_\varphi(\hat{x})) g(\Phi_\varphi(\hat{x}), \Phi_\psi(\hat{y})) \psi_j(\Phi_\psi(\hat{y}))\, d\hat{y}\, d\hat{x}.$$

The Gramian determinants are constant and can be computed directly. Rewriting the integral with the transformed functions

$$\hat{\varphi}_i(\hat{x}) := \varphi_i(\Phi_\varphi(\hat{x})), \qquad \hat{\psi}_j(\hat{y}) := \psi_j(\Phi_\psi(\hat{y})), \qquad \hat{g}(\hat{x},\hat{y}) := g(\Phi_\varphi(\hat{x}), \Phi_\psi(\hat{y}))$$

yields

$$\int_{\widehat{\Delta}} \int_{\widehat{\Delta}} \hat{\varphi}_i(\hat{x}) \hat{g}(\hat{x},\hat{y}) \hat{\psi}_j(\hat{y})\, d\hat{y}\, d\hat{x}. \tag{2}$$

For standard finite element discretizations, the transformed basis functions $\hat{\varphi}_i$ and $\hat{\psi}_j$ are low-order polynomials that can be integrated easily. The transformed kernel function $\hat{g}$, on the other hand, poses a challenge: the integrand behaves differently depending on whether the original triangles Δ_φ and Δ_ψ are identical, share a common edge, a common vertex, or are disjoint. In the last case, the integrand is smooth and we can apply standard quadrature. In the first three cases, the integrand has a singularity on a two-, one- or zero-dimensional sub-manifold of the domain $\widehat{\Delta} \times \widehat{\Delta}$ of integration.

The idea of Sauter-Schwab quadrature [26, 28] is to apply suitable transformations in each of the four cases to eliminate the singularity and obtain integrands that can be handle by standard Gaussian quadrature.

If Δ_φ and Δ_ψ are disjoint, there is no singularity, and we use the Duffy transformation

$$T : [0,1]^2 \to \widehat{\Delta}, \qquad\qquad (s,t) \mapsto (s, st),$$

to write (2) in the form

$$\int_{[0,1]^4} \xi_1 \xi_3\, \hat{\varphi}_i(\xi_1, \xi_1\xi_2)\, \hat{g}(\xi_1, \xi_1\xi_2, \xi_3, \xi_3\xi_4)\, \hat{\psi}_j(\xi_3, \xi_3\xi_4)\, d\xi.$$

This integral can be approximated using standard tensor Gaussian quadrature.

If Δ_φ and Δ_ψ share a common point, but not a common edge, we can choose Φ_φ and Φ_ψ in such a way that $\Phi_\varphi(0) = \Phi_\psi(0)$, i.e., that the origin is mapped to the singularity. Following [28, §5.2.3], the integral (2) can be transformed to

$$\begin{aligned}
&\int_0^1 \xi^3 \int_{[0,1]^3} \eta_2\, \hat{\varphi}_i(\xi, \xi\eta_1)\, \hat{g}(\xi, \xi\eta_1, \xi\eta_2, \xi\eta_2\eta_3)\, \hat{\psi}_j(\xi\eta_2, \xi\eta_2\eta_3)\, d\eta\, d\xi \\
&\quad + \int_0^1 \xi^3 \int_{[0,1]^3} \eta_2\, \hat{\varphi}_i(\xi\eta_2, \xi\eta_2\eta_3)\, \hat{g}(\xi\eta_2, \xi\eta_2\eta_3, \xi, \xi\eta_1)\, \hat{\psi}_j(\xi, \xi\eta_1)\, d\eta\, d\xi,
\end{aligned}$$

and these integrals can again be approximated by tensor Gaussian quadrature.

In the case of a common edge and identical triangles, the procedures in [28, §5.2.2, §5.2.1] can be used to obtain similar equations with five and six four-dimensional integrals, respectively.

We can already see that Sauter-Schwab quadrature is particularly attractive for GPGPUs: the algorithm requires only the coordinates of the vertices of Δ_φ and Δ_ψ and then performs a relatively large number of operations, the evaluation of the kernel function and the basis functions in all quadrature points, so we have a fairly compute-intensive task. It is also clear that we can parallelize both the evaluation of the kernel function for one pair of triangles and the entire quadrature algorithm for entire sets of such pairs.

For the compression of the matrix G, we use the Green cross approximation (GCA) algorithm [5]: we split the matrix into blocks $G|_{t\times s}$ corresponding to subsets $t \subseteq \mathcal{I}$ and $s \subseteq \mathcal{J}$. These subsets are organized in a hierarchical structure that allows us to easily determine the blocks $G|_{t\times s}$ that can be approximated by low-rank matrices in factorized form.

The GCA algorithm constructs the approximations in two steps: first algebraic interpolation operators

$$\mathfrak{I}_t : \mathbb{C}^{\hat{t}} \to \mathcal{V}_t, \qquad \mathfrak{I}_s : \mathbb{C}^{\hat{s}} \to \mathcal{W}_s$$

mapping into low-dimensional spaces $\mathcal{V}_t$ and $\mathcal{W}_s$ are determined using a combination of Green's representation formula, quadrature, and adaptive cross approximation [2]. These operators can be represented in the form

$$\mathfrak{I}_t = V_t P_t, \qquad V_t \in \mathbb{C}^{t\times\tilde{t}}, \qquad \mathfrak{I}_s = W_s P_s, \qquad W_s \in \mathbb{C}^{s\times\tilde{s}},$$

where $\tilde{t} \subseteq t$ and $\tilde{s} \subseteq s$ are small sets of pivot elements chosen by the cross approximation algorithm and

$$P_t u = u|_{\tilde{t}}, \qquad P_s v = v|_{\tilde{s}} \qquad \text{for all } u \in \mathbb{C}^t,\ v \in \mathbb{C}^s$$

are simple restriction operators that pick the components corresponding to pivot elements from larger vectors. Compared to standard Lagrange interpolation, the columns of the matrices V_t and W_s are the counterparts of Lagrange polynomials, while the operators P_t and P_s represent the evaluation of the argument in the interpolation points.

In the second step of the algorithm, we apply the interpolation to all blocks that can be approximated, i.e., we use

$$G|_{t\times s} \approx \mathfrak{I}_t G|_{t\times s} \mathfrak{I}_s^* = V_t G|_{\tilde{t}\times\tilde{s}} W_s^*.$$

Since $\tilde{t}$ and $\tilde{s}$ usually contain a significantly lower number of elements than t and s, this procedure yields a low-rank approximation of $G|_{t\times s}$ in factorized form.

The first step of the compression procedure is adaptive and therefore not particularly well-suited for vector processors, so we leave it to the main processor. This is not a major drawback, since the number of operators $\mathfrak{I}_t$ and $\mathfrak{I}_s$ is significantly lower than the number of blocks.

The second step of the compression procedure requires the computation of the matrices $G|_{\tilde{t}\times\tilde{s}}$, and this task can again be handled by GPGPUs using the procedure described before.

3 Implementation

Although the previously mentioned computations are well suited for GPGPUs, there are a few technical obstacles to consider. In fact, there are two major problems that we have to handle:

- Near-field matrix blocks $G|_{t\times s}$ that cannot be approximated by the GCA technique usually consist of all four types of integrals arising from Sauter-Schwab quadrature.
- Both inadmissible blocks $G|_{t\times s}$ and admissible blocks $G|_{\tilde{t}\times\tilde{s}}$ are fairly small.

The first fact will lead to divergent execution branches when not handled correctly, which would be a major drawback for vectorized execution. When naively executing each block one by one, the second fact could be responsible for a bad utilization of the GPGPU's hardware due to the little amount of work required for a single matrix block.

To address the first issue, we simply keep track of all integrals that appear in the course of the matrix assembly. Depending on the quadrature case, the indices of the triangles and the memory cell the integral belongs to in the matrix will be saved in individual lists for all four cases. As a result this approach will separate all four different types of integrals. If we launch a kernel on the GPGPU for every list separately, no branch divergence effects will occur anymore due to the same execution paths being used for the elements of each list.

A nice side-effect of this approach is that we can control the size of these lists. This also addresses our second issue mentioned before: there is no longer a need for kernel executions with only a few integrals to be computed. Instead, we can increase the maximal size of the lists to a problem-dependent size. Experiments show that the size should be several megabytes to obtain a reasonable performance.

Nevertheless, this algorithmic realization has some drawbacks as well: one has to allocate extra memory to implement the bookkeeping, but as the user can choose the size of the lists and the storage can be re-used after one computation of a list has finished, this is not a major drawback.

On the other hand, after the computations on the GPGPU have finished, the integrals have to be added to their final position in the stiffness matrix G. These memory operations have to be performed by the CPU and in a single-threaded program the GPGPU cannot do any further computations because the CPU is not able to provide work to the GPGPU while copying data. Consequently, we propose to use multiple CPU threads to overcome this problem. (cf. figure 3)

Of course, it would also be beneficial to let the CPU cores compute integrals, since they also contain SIMD units. Here we focus on an approach that has a single thread, in the following called the *master thread*, traverses the block tree and simply fills the lists with work items, while other CPU threads, called *worker threads* have the task to

prepare the data in the lists for the GPGPU computation, launch appropriate kernels on the GPGPU, and distribute the results back to the stiffness matrix. In fact, this is a simple task scheduler following the producer-consumer pattern.

The implementation is based on the hierarchical matrix library *H2Lib*[1] written in the C programming language. For the realization of the task scheduler on the main processor, we employ OpenMP to create the *master* and a various number of *worker* threads. The computations of the GPGPU are carried out using the OpenCL standard which allows us to execute the quadrature code on different hardware platforms that implement OpenCL, not just on GPGPUs. Of course, our approach can also be applied with different standards for multithreading and heterogeneous computing.

We can split the assembly of the stiffness matrix into the following parts:

1. Set up all algebraic interpolation operators $\mathfrak{I}_t$ and $\mathfrak{I}_s$ for all subsets $t \subseteq \mathcal{I}$ and $s \subseteq \mathcal{J}$ on the CPU, since this cannot be properly vectorized.

2. Split the program execution into the execution of a *master thread* and several *worker threads*: Let the master thread traverse the block tree and add each block to the list of disjoint triangles and keep track of its size.

   As soon as the total size of one of these lists exceeds a given limit, mark the list as "ready for execution" and hand it over to the worker threads in the scheduler. This is also depicted in algorithm 1.

3. Once a list has been marked as "ready for execution", a thread of the scheduler can fetch this list (cf. algorithm 2) and process this data package. (see algorithm 3)

4. Once a work item has been processed on the GPU, the results need to be copied to their proper location in the stiffness matrix by the CPU. In case a block could potentially contain non-disjoint pairs of triangles they are now added to the list for their respective case and handled as in (2) which is also displayed in algorithm 3.

Remark 1 *At first sight, treating all matrix entries as if they would belong to the case of distant triangles seems to be counter-intuitive, but it speeds up the setup phase for the GPGPU computation as well as the collection phase.*

It does not significantly slow down the computation on the GPGPU because we treat complete matrix blocks instead of only single entries and the disjoint triangle pairs make up the majority.

We just have to ensure that the results of the non-disjoint pairs overwrite the incorrect coefficients returned by the list of disjoint pairs.

[1] Available at `http://www.h2lib.org`

Algorithm 1 Assembly of a boundary integral operator G. The parameters of all near-field and far-field blocks are initially added to the list of disjoint triangles. The procedure ADD_BLOCK takes care that these lists do not become too big and are queued for execution by the scheduling algorithm once they reach a given size.

```
procedure ADD_BLOCK(b)
    if SIZEOF(ℓ) + SIZEOF(b) < maxsize then
        ADD_TO_LIST(ℓ,b)
    else if sizeof(ℓ) > 0 then
        ENQUEUE_LIST(ℓ)
        ℓ = NEW_LIST(maxsize)
        if SIZEOF(ℓ) + SIZEOF(b) < maxsize then
            ADD_TO_LIST(ℓ,b)
        else
            bs ← SPLIT_BLOCK(b)
            ADD_BLOCK(bs)
        end if
    end if
end procedure

procedure ASSEMBLE_GPU(G)
    for all Matrix blocks b ∈ {G_|t×s, G_|t̂×ŝ} do
        ADD_BLOCK(b)
    end for
end procedure
```

Algorithm 2 Basic sketch of the scheduler algorithm that is used. As long as new work items are available, each thread fetches an item and either executes it on the CPU or on the GPU depending on the affinity of the current item. When running out of work items to be processed, a thread will go to sleep for a short period of time or will terminate itself if an exterior signal was received that indicates that the scheduler should stop.

```
procedure SCHEDULE_TASKS
    while ℓ ← DEQUEUE_LIST do
        if ℓ != empty then
            if AFFINITY(ℓ) == GPU then
                EXECUTE_GPU(ℓ)
            else
                EXECUTE_CPU(ℓ)
            end if
        else
            SLEEP
        end if
    end while
end procedure
```

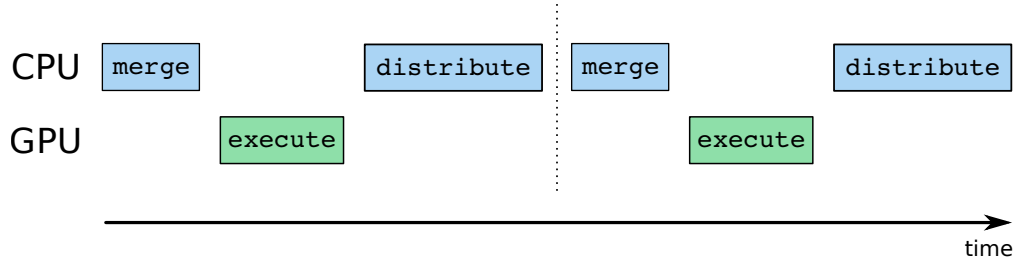

(a) GPU occupancy with a single worker thread

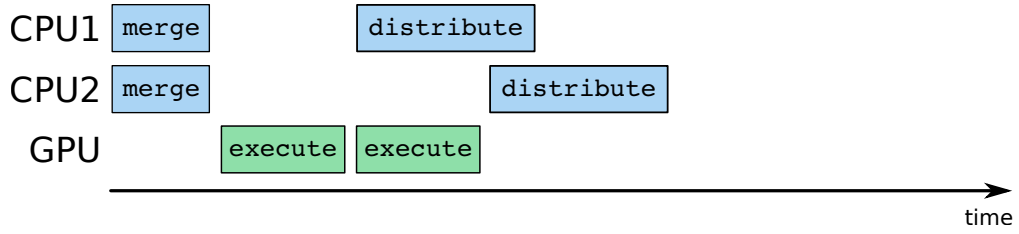

(b) GPU occupancy with two worker threads

Figure 1: Schematic difference of a single worker threads vs. two worker threads supplying a GPU with work items.

Algorithm 3 Result distribution for the case of disjoint triangles. If a block also contains non-disjoint pairs of triangles, they are now added to the list of their respective case and processed in a similar way.

```
procedure DISTRIBUTE_RESULTS_DISJOINT(ℓ, data)
    for all G_|t×s ∈ ℓ do
        COPY_RESULTS(data, G_|t×s)
        if G_|t×s contains non-disjoint pairs of triangles then
            for all (i, j) ∈ G_|t×s do
                v ← NUMBER_OF_COMMON_VERTICES(t, s)
                if v > 0 then
                    ADD_ENTRY_TO_LIST(ℓ_v, (i, j))
                end if
            end for
        end if
    end for
end procedure
```

Algorithm 4 Executing a list of tasks ℓ on the GPU. The parameters from the individual tasks needs to be prepared into big chunks of parameters. Later-on a kernel on the GPU is invoked and results will be distributed to their proper location in main memory.

```
procedure EXECUTE_GPU(ℓ)
    data ← MERGE_DATA(ℓ)
    INVOKE_GPU_KERNEL(data)
    DISTRIBUTE_RESULTS(ℓ, data)
    CLEANUP(data)
end procedure
```

4 Numerical experiments

In this part we want to demonstrate the performance of our GPGPU algorithm for setting up boundary integral operators arising in the boundary element method.

For the following experiments we consider the solution of an interior Dirichlet problem for the Laplace equation in 3D. Let f be a harmonic function in a bounded domain $\Omega \subset \mathbb{R}^3$ and assume that the Dirichlet values $f|_{\partial\Omega}$ are given. In the context of the boundary element method, this translates into the boundary integral equation

$$\int_{\partial\Omega} g(x,y)\frac{\partial f}{\partial n}(y)\mathrm{d}y = \frac{1}{2}f(x) + \int_{\partial\Omega}\frac{\partial g}{\partial n(y)}(x,y)f(y)\mathrm{d}y \quad \text{for all } x \in \partial\Omega, \tag{3}$$

where we are looking for the Neumann values $\frac{\partial f}{\partial n}$. Here

$$g(x,y) := \frac{1}{4\pi\|x-y\|}$$

denotes the fundamental solution of the Laplace operator in 3D.

A Galerkin discretization of this equation with piece-wise constant basis functions $(\varphi_i)_{i\in\mathcal{I}}, (\psi_j)_{j\in\mathcal{J}}$ leads to a linear system

$$V\alpha = \left(\frac{1}{2}I + K\right)\beta,$$

where the coefficients of the matrices V and K are given by equation (1) with the kernel functions g and $\frac{\partial g}{\partial n(y)}$, respectively. The values of f and $\frac{\partial f}{\partial n}$ are represented in the bases as $f = \sum_{i\in\mathcal{I}}\alpha_i\varphi_i$ and $\frac{\partial f}{\partial n} = \sum_{j\in\mathcal{J}}\beta_j\psi_J$, respectively. Since the matrix V is symmetric and positive definite, we can employ a conjugate gradient method [21] to solve the linear system.

In fact the most time-consuming part of the whole computation is the setup of the matrices V and K, which is carried out mostly by the GPGPU. In order to verify the correctness of our proposed method, we solve the system with the following three harmonic functions:

$$f_1(x) = x_1^2 - x_3^2, \quad f_2(x) = g(x,(1.2,1.2,1.2)), \quad f_3(x) = g(x,(1.0,0.25,1.0)).$$

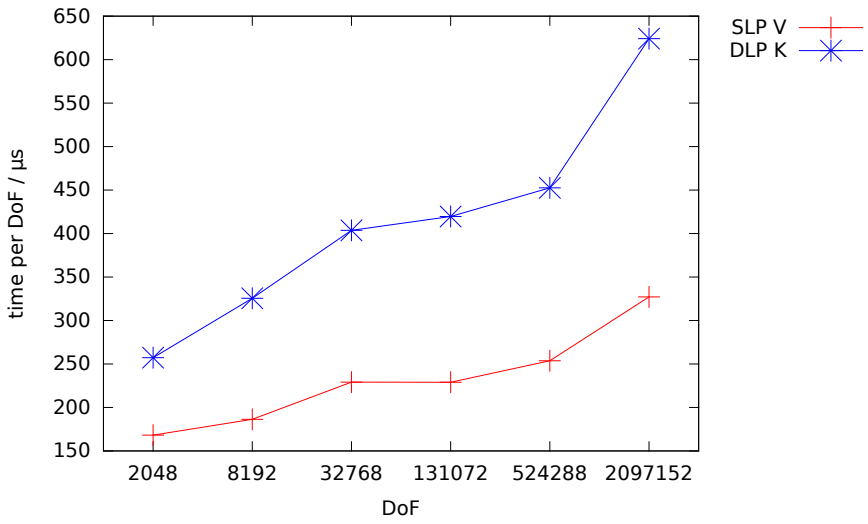

Figure 2: Setup time of the boundary integral operators on different resolutions of the unit sphere. Only CPU threads are employed to calculate the integrals.

As a simple test domain, we choose Ω as the unit sphere. Due to the symmetry of V, we only have to compute its lower triangular part. Therefore the setup times for K are roughly twice those for the matrix V, even more since the kernel function of K requires more operations than the one for V.

CPU algorithm To have a reference point for our further experiments, we first present some results computed entirely on the CPU. Instead of traversing the block tree and keeping account of the the arising matrix sub-blocks to be computed, the CPU algorithm directly computes these matrix entries. The algorithm was executed on a Intel Core i7-3820 processor with 4 cores and hyper-threading enabled, which allows us to use up to 8 threads on the CPU. Consequently we want to use all the 8 threads to execute the algorithm on the CPU in order to keep pace with the GPUs later on. Explicit vectorization for the CPU cores is not enabled, although the CPU has AVX vector instructions at its disposal. The results can be seen in figure 2. The timings do not include the setup time for the algebraic interpolation operators, since these are always precomputed by a CPU algorithm. In this and all later experiments we scale the runtime by the number of degrees of freedom in order to compare the values across a wide range of resolutions of the geometry. We can see in that figure that the runtime on the main processor is between 150 and 650 μs for both SLP operator V and DLP operator K for all resolutions. Although the GCA algorithm can compute an approximation to both matrices in linear time with respect to the number of unknowns, we need to lower the error tolerances in each step of the geometrical refinement to reach the discretization error in our solutions.

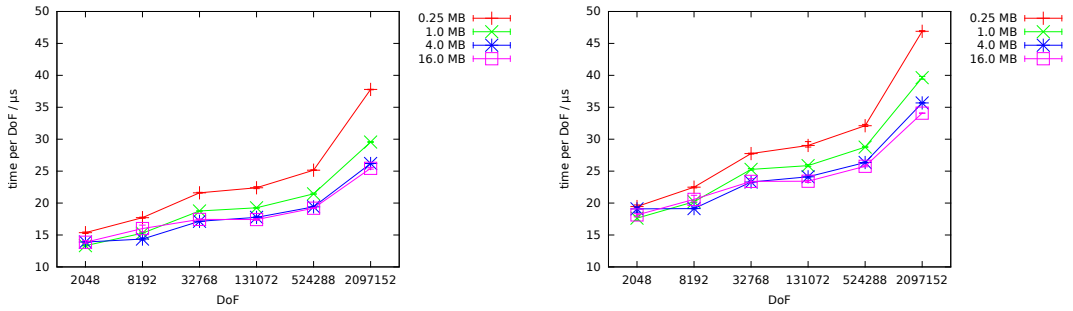

(a) Runtime for V on GeForce GTX 580

(b) Runtime for V on GeForce GTX 680

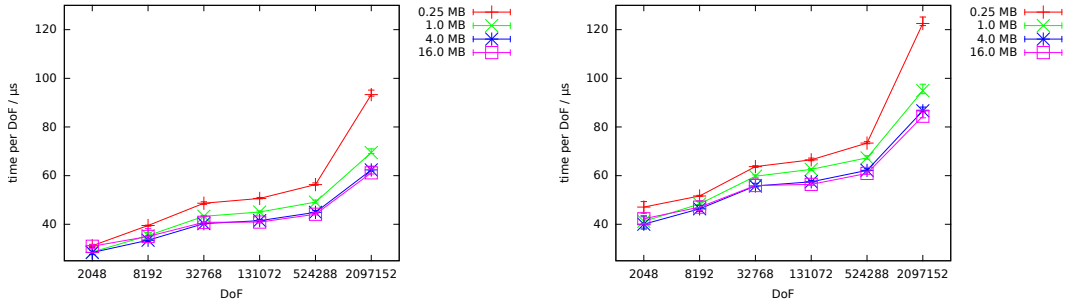

(c) Runtime for K on GeForce GTX 580

(d) Runtime for K on GeForce GTX 680

Figure 3: Setup time of the boundary integral operators for varying list sizes on different resolutions of the unit sphere. An NVIDIA GeForce GTX 580 was used to compute the SLP operator in 3(a) and the DLP operator in 3(c). Likewise an NVIDIA GeForce GTX 680 was used to compute the SLP operator in 3(b) and the DLP operator in 3(d).

List size Now we consider the performance of our GPGPU algorithm, particularly in relation to the parameters of the approach, namely the size of lists, as well as the number of worker threads that supply the GPUs. Regarding the size of the lists it seems obvious that a smaller size automatically corresponds to more kernel executions on the GPU. Since each call of a GPU kernel involves some overhead, it is not surprising that smaller list sizes lead to longer execution times of our algorithm. As shown by figure 3, we solved the system for different resolutions of the sphere, up to 2.1 million degrees of freedom, and looked again at the setup time of the matrices with varying size of the lists. The parameters of the GCA approximation scheme are chosen as in [5] to satisfy $\mathcal{O}(h)$ convergence on the Neumann values close to the discretization error.

Due to non-deterministic thread scheduling of the operating system and the thread scheduler on the graphics card, we always run the computations five times on the GPU and note the minimal, maximal and average runtime of the algorithm. This is depicted by the error-bars in figure 3 and the following ones. One can see that the deviation

is larger for very small list sizes and negligible for larger ones. The main conclusion drawn from figure 3 is that the list size should be chosen somewhere around 4–16 MB to minimize the runtime. However this is not the best choice in all cases. In general there is no best choice for all degrees of freedom for this parameter, but 8 MB seems to be a good empirical value.

Remark 2 *At first sight the results from figure 3 seem to be confusing because the GeForce GTX 680 achieves a lower performance than its architectural predecessor, the GeForce GTX 580. But in fact, this can be explained by changes in the GPU's internal structure, especially for double precision computations. The theoretical peak performance of the GeForce GTX 680 is given by 128.8 GFlops and the peak performance of the GeForce GTX 580 is given as 197.6 GFlops in double precision. Therefore we can expect the GeForce GTX 580 to be faster than the GeForce GTX 680 by a factor of 1.5 which is confirmed by our results.*

Worker threads Since there is always a portion of time in between two successive kernel calls on the GPU where the CPU has to perform some data management operations, the utilization of the GPU will not be near its maximum. Therefore our implementation of the algorithm has a parameter controlling the number of CPU worker threads spawned per GPU to supply the GPU with work. We now investigate the effect of the number of worker threads on the overall performance of our approach.

In figure 4 we can see a significant reduction of the runtime when two or more worker threads are employed per GPU. It is not worth while to use more than two threads per GPU. Indeed using more than two threads might have a negative impact on the performance if not enough hardware threads are available to supply the GPUs with sufficient work or the memory bandwidth of the CPU might be exhausted at some point. Also one can observe that the deviation in runtime increases for smaller resolutions of the geometry and decreases for bigger resolutions as more worker threads are employed. This might be due to the non-deterministic behaviour of CPU thread scheduling by the operating system. Additionally, for very small number of unknowns using more worker threads can also influence the runtime in a negative way because there is not enough work to distribute and some threads might be idle.

Multi-GPU setup In the design of our algorithm, we took into account that a system can have more than just one GPGPU available. A single worker thread is always directly connected with a single GPU and cannot supply work to a different GPU. With this concept it is very comfortable to employ all GPUs in a system to work on the same problem. On the other hand, all GPUs are connected by the same PCIe bus system in typical computer systems, and it is not quite clear if the total amount of data that is being transferred over this bus can be handled without any delay. We investigate this behaviour in the next experiment, where we employ both a GeForce GTX 580 and GeForce GTX 680 at the same time. In this context we are interested in how the runtime of our algorithm scales when using both cards simultaneously. figure 5 illustrates the effect of the multi-GPU setup.

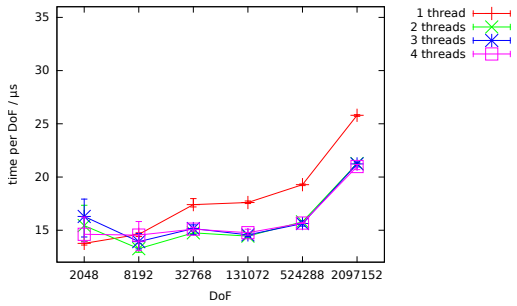

(a) Runtime for V on GeForce GTX 580

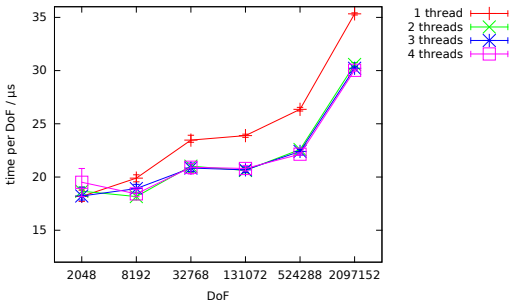

(b) Runtime for V on GeForce GTX 680

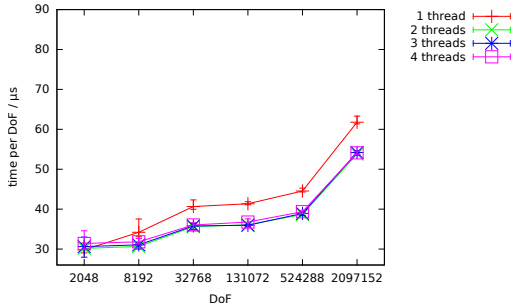

(c) Runtime for K on GeForce GTX 580

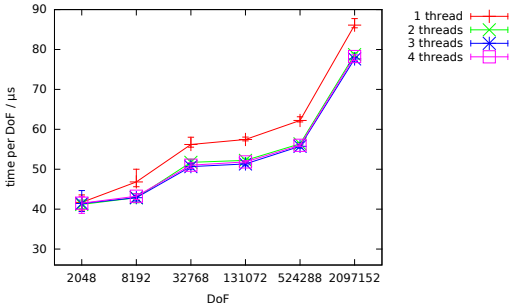

(d) Runtime for K on GeForce GTX 680

Figure 4: Setup time of the boundary integral operators for varying number of worker threads per GPU on different resolutions of the unit sphere. The list size was chosen constant as 8 MB. An NVIDIA GeForce GTX 580 was used to compute the SLP operator in 4(a) and the DLP operator in 4(c). Likewise an NVIDIA GeForce GTX 680 was used to compute the SLP operator in 4(b) and the DLP operator in 4(d).

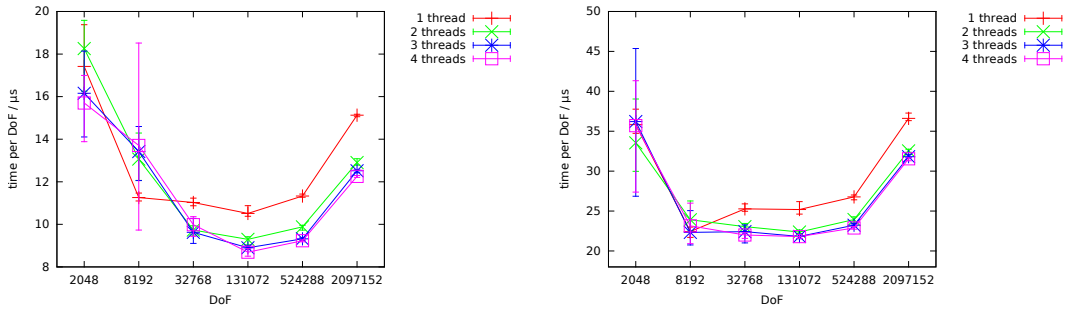

(a) Runtime for V on GeForce GTX 580 and 680 (b) Runtime for K on GeForce GTX 580 and 680

Figure 5: Setup time of the boundary integral operators using both GeForce GTX 580 and GeForce GTX 680 to compute the integrals. The list size is kept constant at 8 MB and two worker threads are being employed per GPU.

Table 1: Runtimes for CPU and GPGPU algorithms to setup SLP operator V. List size is kept constant at 8 MB and 2 worker threads are employed.

DoFs	$\mathcal{I}_t, \mathcal{I}_s$	CPU	GTX 580	GTX 680	GTX 580 + 680
2048	<0.1 s	0.3 s	<0.1 s	<0.1 s	<0.1 s
8192	0.1 s	1.5 s	0.1 s	0.2 s	0.1 s
32768	0.4 s	7.5 s	0.5 s	0.7 s	0.3 s
131072	1.5 s	30.0 s	1.9 s	2.7 s	1.2 s
524288	6.2 s	133.0 s	8.3 s	11.8 s	5.2 s
2097152	62.0 s	686.2 s	44.6 s	64.1 s	27.1 s

At least for this two-GPU setup, our approach seems to scale well. According to our previous remark about the peak performances of both GPUs we can only expect the runtime for the multi GPU setup of our algorithm to be smaller by a factor of 1.7 compared to the runtime of the GeForce GTX 580 or smaller by a factor of 2.5 compared to the runtime of the GeForce GTX 680. Indeed our experiments show similar factors, and we can say that our scheduler handles the two-GPU setup quite well.

The scaling behaviour of our algorithm on more than two GPUs is an interesting topic and will be part of future research.

Tables 1 and 2 list the absolute computing times for our algorithm and also include setup time for the algebraic interpolation operators for the sake of completeness, which are entirely computed on the CPU. One can observe that the time that is necessary to set up the interpolation operators is roughly the same as the time required by the setup of the near-field and coupling matrices on both GPUs.

We could speed up computation of the interpolation operators as well, either by using SIMD instructions on the CPU or with the help of GPGPUs.

More important is the time needed to solve the three linear systems. We used a

Table 2: Runtimes for CPU and GPGPU algorithms to setup DLP operator K. List size is kept constant at 8 MB and 2 worker threads are employed.

DoFs	$\mathcal{I}_t, \mathcal{I}_s$	CPU	GTX 580	GTX 680	GTX 580 + 680
2048	<0.1 s	0.5 s	<0.1 s	<0.1 s	<0.1 s
8192	0.1 s	2.7 s	0.3 s	0.4 s	0.2 s
32768	0.5 s	13.2 s	1.2 s	1.7 s	0.8 s
131072	1.7 s	55.0 s	4.7 s	6.8 s	2.9 s
524288	6.7 s	237.3 s	20.3 s	29.5 s	12.5 s
2097152	70.2 s	1309.0 s	112.8 s	165.1 s	68.2 s

Table 3: Runtimes for CPU and GPGPU algorithms to setup SLP V and DLP operator K for basic quadrature rule using 4(6) points. List size is kept constant at 8 MB and 2 worker threads are employed. For the last line, the GCA parameter $\delta = 2$ was chosen instead of $\delta = 1$ to make the problem fit into memory.

DoFs	CPU SLP	GTX 580 + 680 SLP	CPU DLP	GTX 580 + 680 DLP
2048	1.0 s	<0.1 s	1.5 s	0.1 s
8192	4.8 s	0.2 s	7.7 s	0.4 s
32768	21.6 s	0.7 s	39.0 s	1.8 s
131072	86.8 s	2.8 s	162.1 s	7.3 s
524288	368.5 s	12.0 s	702.0 s	31.3 s
2097152	1439.3 s	48.3 s	2936.2 s	131.9 s

simple conjugate gradient method which needed 260 steps on the finest mesh and took about 1400 seconds per right-hand side. At this point a good preconditioner would be very desirable. The *H2Lib* package offers $\mathcal{H}^2$-Cholesky and $\mathcal{H}^2$-LU decompositions, but unfortunately these are not yet parallelized. Hence a comparison with a sequential algorithm for constructing a preconditioner would not be fair.

Higher quadrature orders All experiments so far have been conducted with a basic one dimensional Gaussian quadrature formula with 3 points for Duffy transformations and with 5 points for the Sauter-Schwab quadrature formulas. In some applications this quadrature order might not be sufficient if the geometries become less smooth. Therefore we also present some results for a basic Gaussian quadrature formula with 4(6) quadrature points in Table 3. Since the unit sphere is a very smooth geometry, such a high quadrature order is not really necessary, but we can see that the GPUs benefit from the higher order, since their efficiency increases with the amount of work required for a single integral. Hence the speedup factor compared to the CPU implementation rises from 25.3 to 29.8 for the single layer potential and from 19.2 to 22.3 for the double layer potential.

Helmholtz equation As a final example we consider a slightly different problem setting: the Helmholtz equation in its BEM formulation. In this case we would like to solve the exterior Dirichlet problem, given by

$$\begin{aligned} \Delta u(x) + \kappa^2 u(x) &= 0, & x &\in \mathbb{R}^3 \setminus \bar{\Omega} \\ u(x) &= f(x), & x &\in \partial\Omega, \\ \left| \left\langle \nabla u(x), \frac{x}{\|x\|} \right\rangle - i\kappa u(x) \right| &\in \mathcal{O}\left(\frac{1}{\|x\|^2}\right) & \|x\| &\to \infty. \end{aligned}$$

Here κ denotes the wave number. It is well known, that if κ is related to an eigenvalue of the interior Neumann problem, then the standard Dirichlet-to-Neumann mapping (3) is not applicable, due to the single layer potential operator not being coercive anymore. Well-known proposals to fix this problem are the Brakhage-Werner approach [10] and the Burton-Miller approach [11]. We choose the former and reformulate the problem as

$$\int_{\partial\Omega} \frac{\partial g}{\partial n(y)}(x,y) w(y) \mathrm{d}y - i\eta \int_{\partial\Omega} g(x,y) w(y) \mathrm{d}y = f(x) \quad \text{for all } x \in \mathbb{R}^3 \setminus \bar{\Omega},$$

which has to be solved for the unknown density w. To ensure unique solvability, we have to choose $\eta > 0$. Applying trace operators yields a boundary integral equation which we can solve with the same techniques as (3).

Again we define single and double layer potential operators V and K using the fundamental solution

$$g(x,y) := \frac{e^{i\kappa\|x-y\|}}{\|x-y\|}$$

of the Helmholtz equation. Now we want to demonstrate the computing power of the two GPUs combined against all CPU cores. In the same way as we did for Laplace's equation, we used the unit sphere as geometry and computed the solution up to discretization error for different resolutions of the sphere, but with a constant wave-number $\kappa = 3$. As Dirichlet data we choose the function

$$f(x) = g(x, (0.0, 0.0, 0.2)).$$

Since we do not have an analytic solution for w at our disposal, we evaluate the solution u of the Helmholtz equation by the Brakhage-Werner formulation and compare it to the analytically solution of the problem in the exterior of the domain, given by the function f itself. The results can be seen in Table 4.

5 Conclusions

Matrices that arise from the Galerkin discretization of boundary integral equations are densely populated and are therefore very compute-intensive. We have proposed an implementation of the Green cross approximation algorithm on GPGPUs and have shown that this approach is very fast compared to pure CPU code, allowing us to set up very large problems quickly.

Table 4: Runtimes for CPU and GPGPU algorithms to setup SLP V and DLP operator K for the Helmholtz equation. List size is kept constant at 8 MB and 2 worker threads are employed.

DoFs	CPU SLP	GTX 580 + 680 SLP	CPU DLP	GTX 580 + 680 DLP
2048	2.3 s	0.1 s	2.7 s	0.1 s
8192	9.6 s	0.3 s	11.8 s	0.4 s
32768	42.7 s	1.3 s	55.2 s	1.8 s
131072	205.8 s	6.7 s	263.4 s	8.9 s
524288	843.6 s	27.9 s	1050.1 s	36.9 s

We have also shown that both, the GCA method and our implementation on the GPU, work not only for the Laplace equation, but also for the significantly more challenging Helmholtz equation.

Our method is not limited to computation on graphic cards, but can utilize general computing hardware supporting the OpenCL standard. Especially the use of more that just one accelerator card makes this approach very interesting for high performance computing applications because it might be scaled to several hundreds of GPUs with some modifications.

Until now, the setup phase of the boundary integral matrices has dominated the computing time of boundary element methods. Using our new CPU/GPU approach, the time required for solving the linear system has become the limiting factor. This challenge could be approach, e.g., by parallelizing $\mathcal{H}$- and $\mathcal{H}^2$-matrix preconditioners [14, 9, 22].

References

[1] C. R. Anderson. An implementation of the fast multipole method without multipoles. *SIAM J. Sci. Stat. Comp.*, 13:923–947, 1992.

[2] M. Bebendorf and S. Rjasanow. Adaptive Low-Rank Approximation of Collocation Matrices. *Computing*, 70(1):1–24, 2003.

[3] M. Bebendorf and R. Venn. Constructing nested bases approximations from the entries of non-local operators. *Num. Math.*, 121(4):609–635, 2012.

[4] S. Börm. *Efficient Numerical Methods for Non-local Operators: $\mathcal{H}^2$-Matrix Compression, Algorithms and Analysis*, volume 14 of *EMS Tracts in Mathematics*. EMS, Zürich, 2010.

[5] S. Börm and S. Christophersen. Approximation of integral operators by Green quadrature and nested cross approximation. *Num. Math.*, pages 1–34, 2015. to appear.

[6] S. Börm and J. Gördes. Low-rank approximation of integral operators by using the Green formula and quadrature. *Numerical Algorithms*, 64(3):567–592, 2013.

[7] S. Börm and L. Grasedyck. Hybrid cross approximation of integral operators. *Numer. Math.*, 101:221–249, 2005.

[8] S. Börm and W. Hackbusch. Data-sparse approximation by adaptive $\mathcal{H}^2$-matrices. *Computing*, 69:1–35, 2002.

[9] S. Börm and K. Reimer. Efficient arithmetic operations for rank-structured matrices based on hierarchical low-rank updates. *Comp. Vis. Sci.*, 16(6):247–258, 2015. `http://arxiv.org/abs/1402.5056`.

[10] H. Brakhage and P. Werner. Über das dirichletsche Außenraumproblem für die helmholtzsche schwingungsgleichung. *Archiv der Mathematik*, 16(1):325–329, 1965.

[11] A. J. Burton and G. F. Miller. The application of integral equation methods to the numerical solution of some exterior boundary-value problems. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 323(1553):pp. 201–210, 1971.

[12] S. Erichsen and S. A. Sauter. Efficient automatic quadrature in 3-d Galerkin BEM. *Comput. Meth. Appl. Mech. Eng.*, 157:215–224, 1998.

[13] Z. Gimbutas and V. Rokhlin. A generalized fast multipole method for nonoscillatory kernels. *SIAM J. Sci. Comput.*, 24(3):796–817, 2002.

[14] L. Grasedyck and W. Hackbusch. Construction and arithmetics of $\mathcal{H}$-matrices. *Computing*, 70:295–334, 2003.

[15] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comp. Phys.*, 73:325–348, 1987.

[16] W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices. *Computing*, 62:89–108, 1999.

[17] W. Hackbusch and B. N. Khoromskij. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part II: Application to multi-dimensional problems. *Computing*, 64:21–47, 2000.

[18] W. Hackbusch, B. N. Khoromskij, and S. A. Sauter. On $\mathcal{H}^2$-matrices. In H. Bungartz, R. Hoppe, and C. Zenger, editors, *Lectures on Applied Mathematics*, pages 9–29, Berlin, 2000. Springer-Verlag.

[19] W. Hackbusch and Z. P. Nowak. On the fast matrix multiplication in the boundary element method by panel clustering. *Numer. Math.*, 54:463–491, 1989.

[20] S. Hamada. Performance comparison of three types of gpu-accelerated indirect boundary element method for voxel model analysis. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 26(4):337–354, 2013.

[21] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Res. Nat. B. Stand.*, 49(6), 1952.

[22] R. Kriemann. $\mathcal{H}$-LU factorization on many-core systems. *Comp. Vis. Sci.*, 16(3):105–117, 2013.

[23] J. Labaki, L. O. S. Ferreira, and E. Mesquita. Constant boundary elements on graphics hardware: a gpu-cpu complementary implementation. *Journal of the Brazilian Society of Mechanical Sciences and Engineering*, 33:475 – 482, 12 2011.

[24] G. Of, O. Steinbach, and W. L. Wendland. The fast multipole method for the symmetric boundary integral formulation. *IMA J. Numer. Anal.*, 26:272–296, 2006.

[25] V. Rokhlin. Rapid solution of integral equations of classical potential theory. *J. Comp. Phys.*, 60:187–207, 1985.

[26] S. A. Sauter. *Über die effiziente Verwendung des Galerkin-Verfahrens zur Lösung Fredholmscher Integralgleichungen*. Doctoral thesis, Christian-Albrechts-Universität Kiel, 1992.

[27] S. A. Sauter. Variable order panel clustering (extended version). Preprint 52/1999, Max-Planck-Institut für Mathematik, Leipzig, Germany, 1999.

[28] S. A. Sauter and C. Schwab. *Boundary Element Methods*. Springer, Berlin, Heidelberg, 2011.

[29] C. Schwab and W. L. Wendland. Numerical evaluation of singular and finite-part integrals on curved surfaces using symbolic manipulation. *Computing*, 49(3):279–301, 1992.

[30] T. Takahashi and T. Hamada. Gpu-accelerated boundary element method for helmholtz' equation in three dimensions. *International Journal for Numerical Methods in Engineering*, 80(10):1295–1321, 2009.

[31] Y. Wang, Q. Wang, X. Deng, Z. Xia, J. Yan, and H. Xu. Graphics processing unit (gpu) accelerated fast multipole {BEM} with level-skip {M2L} for 3d elasticity problems. *Advances in Engineering Software*, 82:105 – 118, 2015.

[32] L. Ying, G. Biros, and D. Zorin. A kernel-independent adaptive fast multipole algorithm in two and three dimensions. *J. Comp. Phys.*, 196(2):591–626, 2004.